Plants and Productivity in International Trade[1]

Andrew B. Bernard[***], Jonathan Eaton[*], J. Bradford Jensen[**], and Samuel Kortum[*]

[*]Boston University and NBER,
[**]Center for Economic Studies, Bureau of the Census, and University of Maryland
[***]Tuck School of Business, Dartmouth College, and NBER

April 2000

[1]Send correspondence to Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA; jeaton@bu.edu, kortum@bu.edu. We thank Daniel Ackerberg, Eli Berman, Zvi Eckstein, Simon Gilchrist, Nina Pavcnik, and Robert Staiger for helpful comments. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation. The paper was completed while Kortum was a National Fellow at the NBER. Any opinions expressed are those of the authors and not those of the Bureau of the Census, the NSF, or the NBER.

Abstract

We reconcile international trade theory with findings of enormous plant-level heterogeneity in exporting and productivity. Our model extends basic Ricardian theory to accommodate many countries, geographic barriers, and imperfect competition. Fitting the model to bilateral trade among the United States and its 46 major trade partners, we see how well it can explain basic facts about U.S. plants: (i) productivity dispersion, (ii) the productivity advantage of exporters, (iii) the small fraction who export, (iv) the small fraction of revenues from exporting among those that do, and (v) the much larger size of exporters. We pick up all these basic qualitative features, and go quite far in matching them quantitatively. We examine counterfactuals to assess the impact of various global shifts on productivity, plant entry and exit, and labor turnover in U.S. manufacturing.

Key words: International trade, exporting, productivity, heterogeneity, Census of Manufactures

JEL classification: F11; F17; O33

1 Introduction

A new empirical literature has emerged that examines international trade at the level of individual producers. Bernard and Jensen (1995, 1999a), Clerides, Lach, and Tybout (1998), and Aw, Chung, and Roberts (1998), among others, have uncovered stylized facts about the behavior and relative performance of exporting firms and plants which hold consistently across a number of countries: Most strikingly, exporters are in the minority; they tend to be more productive and larger; yet they usually export only a small fraction of their output. This heterogeneity of performance diminishes only modestly when attention is restricted to producers within a given industry.

International trade theory has not had much to say about these producer-level facts, and in many cases is inconsistent with them. To the extent that empirical implications have been of concern, trade theory has been aimed at understanding aggregate evidence on such topics as the factor content of trade and industry specialization. To understand the effects of trade on micro issues such as plant closings, however, we need a theory that recognizes differences among individual producers within an industry. Moreover, as we elaborate below, such a theory is needed to understand the implications of trade for such aggregate magnitudes as worker productivity.

Our purpose here is to develop a model of international trade that comes to grips with what goes on at the producer level. Such a model requires three crucial elements. First, we need to acknowledge the heterogeneity of plants. To do so we introduce Ricardian differences in technological efficiency across producers and countries. Second, we need to explain the coexistence, even within the same industry, of exporters and purely domestic producers. To capture this fact we introduce costs to exporting through a standard "iceberg" assumption (export costs to a given destination are proportional to production costs). Third, in order for differences in technological efficiency not to be fully absorbed by differences in output prices (thus eliminating differences in measured productivity across plants), we need imperfect competition with variable markups. We take the simplest route of introducing Bertrand competition into the Ricardian framework with a given set of goods.[1]

A novel feature of our theory is to link a plant's underlying technological efficiency to its productivity as normally measured (typically as value added per worker). In fact, as long as all producers in a country employ inputs in the same proportion at the same cost, under perfect competition (or, for that matter, monopolistic competition with a common markup) they would all appear equally productive in terms of value added per worker, in spite of any efficiency differences. In our model, differences in value added per unit of input reflect different markups of price over cost which emerge through Bertrand competition. In the absence of any link between efficiency and markups, measured productivity would say nothing about underlying efficiency. It turns out, however, that the theory implies that producers who are more efficient also tend to have a greater cost advantage over their closest competition, and are thus able to set higher markups. Hence measured productivity provides a signal, albeit a noisy one, about underlying technical efficiency.[2] The link between underlying efficiency and productivity, on the one hand, and underlying efficiency and selection into export markets, on the other, leads our model to predict that exporting plants are more productive.[3]

A novel feature of our empirical approach is to connect the micro and macro level data.

[1]As in Eaton and Kortum (2000), specialization emerges endogenously through the exploitation of comparative advantage. An alternative model that also allows for heterogeneity and geographic barriers of the iceberg variety is Krugman's (1979) extension to international trade of the Dixit-Stiglitz (1977) model of monopolistic competition. But this approach delivers the counterfactual implication that every producer exports everywhere. In contrast, in our model a plant exports only when its cost advantage over its competitors around the world overcomes geographic barriers. Other attempts to explain producer heterogeneity in export performance rely on a *fixed* cost of exporting (see, e.g., Roberts and Tybout 1997 and Melitz 1999). The problem here is that a producer would either export nothing or else sell to different countries of the world in proportion to their market sizes. This second implication belies the very small share of exports in the revenues of most exporters.

[2]An extensive literature compares productivity levels across plants. See, e.g., Baily, Hulten, and Campbell (1992), Bartelsman and Dhrymes (1992), and Olley and Pakes (1996). In making such comparisons, it is typically assumed that the plants in question produce a homogeneous output. Our framework shows how such comparisons make sense even when outputs are heterogeneous.

[3]Clerides, Lach, and Tybout (1998) and Bernard and Jensen (1999a) find strong empirical support for this selection mechanism (and little or no empirical support for learning by exporting) in explaining why exporters are more productive than nonexporting plants.

Aggregate production and bilateral trade volumes around the world provide all we need to know about parameters governing geographic barriers, aggregate technology differences, and differences in input costs. The two remaining parameters relate to the heterogeneity of goods in production and in consumption. Reasonable values for these two parameters bring us quite close to fitting the various micro facts as they apply to U.S. manufacturing plants. Hence the framework serves as a bridge between what we know about global trade flows and what we have learned about plant-level export behavior.

Since the model comes to terms with plant-level facts quite well, we go on to ask what changes in the global economy mean for plant entry, exit, and exporting as well as overall productivity and employment in manufacturing. We look at three scenarios.

We first consider the effects of "globalization" in the form of a 5 percent drop in all geographic barriers between countries (resulting in nearly a 40 percent rise in world trade). We find that this move kills off nearly 9 percent of U.S. plants. But among the survivors, one in seven of the plants that had previously sold only to the domestic market starts exporting. Since globalization provides the survivors larger markets, and since the survivors were larger to begin with, the decline in manufacturing employment is less than 3 percent.

We then move in the opposite direction to autarky (raising geographic barriers to eliminate all trade). The number of active U.S. plants rises by 17 percent. But since plants that were exporting lose their overseas markets, nearly as many jobs are destroyed (11 percent of initial employment) as are created (12 percent of initial employment). Reallocation of employment away from exporters to less productive entrants lowers productivity by around 4 percent.

Our final experiment is a decline in U.S. "competitiveness" in the form of an exogenous 10 percent increase in the U.S. relative wage. The number of manufacturing plants falls by 8 percent and manufacturing employment falls by 18 percent as plants substitute cheaper imported intermediates for labor.

Our analysis thus captures how, even in a relatively closed market such as the United

States, changes in the global economy can substantially reshuffle production. This reshuffling in turn can have important implications for overall manufacturing productivity.[4]

We proceed as follows. Section 2 discusses the plant-level facts we seek to explain. In Section 3 we present the theory behind our qualitative explanations, derived in Section 4, for what goes on at the plant level. Section 5 goes on to compare the model's quantitative implications with the plant-level statistics. Section 6 completes the general equilibrium specification of the model required to undertake the counterfactual experiments reported in Section 7. Section 8 concludes.

2 Exporter Facts

Before turning to the theory, we take a closer look at the plant-level statistics about U.S. exporters that our model seeks to explain. These statistics, shown in Table 1, are all calculated from the 1992 U.S. Census of Manufactures (see Appendix B.2).

We first look at the prevalence of exporting among U.S. manufacturing plants. At one extreme, each plant could export the same share of its total output. At the other, a few giant plants would account for all exports. In fact, of the roughly 200,000 plants in the Census, only 21 percent report exporting anything.[5]

While previous work has sought to link trade orientation with industry, it turns out that exporting producers are quite spread out across industries. Figure 1 plots the distribution of industry export intensity: Each of the 458 4-digit manufacturing industries is placed in one of 10 bins according to the percentage of plants in the industry that export. In two-thirds of the industries, the fraction of plants that export lies between 10 and 50 percent. Hence

[4]The results of our counterfactual experiments accord well with findings in the literature. Bernard and Jensen (1999b) find productivity gains driven by reallocation among U.S. producers as exporting has increased. Campa and Goldberg (1995) show that imported intermediates are an important link between U.S. producers and the rest of the world. Gourinchas (1999) estimates that changes in the U.S. real exchange rate lead to increased churning in the labor market. Head and Reis (1998) document the substantial exit and reallocation of production among Canadian producers following tariff reductions under the Free Trade Agreement.

[5]Note that the U.S. export share is even lower. For 1992, the OECD (1995) reports that the U.S. manufacturing sector exported about 13 percent of its gross production.

knowing what industry a plant belongs to leaves substantial uncertainty about whether it exports. Industry has less to do with exporting than standard trade models might suggest.

Not only are plants heterogeneous in whether they export, they also differ substantially in measured productivity. Figure 2A plots the distribution across plants of value added per worker (segregating exporters and nonexporters) relative to the overall mean. A substantial number of plants have productivity either less than a fourth or more than four times the average. Again, a plant's industry is a weak predictor of its performance: Figure 2B provides the same distribution only normalizing each plant's productivity by mean productivity in its 4-digit industry. Controlling for industry only marginally tightens the productivity distribution.

While there is substantial heterogeneity in both productivity and export performance, even within industries, Figure 2A brings out the striking association between the two. The exporters' productivity distribution is a substantial shift to the right of the nonexporters' distribution. Figure 2B shows that this association survives even when looking within 4-digit industries. As shown in Table 1, exporters have a 33 percent advantage in labor productivity overall, and a 15 percent advantage relative to nonexporters within the same 4-digit industry. Accounting for differences in capital intensity across plants within industries, the total factor productivity advantage of exporters is 10 percent.

While differences across industries certainly appear in the data, what is surprising is how little industry explains about exporting and productivity. Hence a satisfactory explanation of plant level behavior must go beyond the industry dimension. We consequently pursue an explanation of these facts that, as a first approximation, bypasses industries and goes directly to the plant level.

Table 1 also reports the importance of export markets for the plants that do export. Surprisingly, the vast majority of exporters export less than 10 percent of what they produce. Less than 5 percent of the exporting plants (which also account for about 5 percent of exporters' total output) export more than 50 percent of their production. Even for the minority of plants that do export, domestic sales dominate.

How is it possible for such a small fraction of plants, exporting such a small fraction of what they produce, to account for total exports? An answer is that exporters are much larger. They are almost 5 times the size of nonexporters on average, even when export revenues are excluded from the calculation. While only 21 percent of manufacturing plants report that they export, these plants account for 60 percent of the output of U.S. manufacturing.

An important caveat in considering any of these statistics is that U.S. manufacturing plants as a whole report exports that sum to just over 60 per cent of total U.S. exports of manufactures reported by the OECD. (See Bernard and Jensen (1995) for a discussion of this problem.) We discuss how undercounting could affect our results below.

3 The Basic Model

Our model combines imperfect competition with the Ricardian theory of comparative advantage based on technology differences. We start with Eaton and Kortum (2000) (Henceforth EK), which itself extends the Ricardian model of Dornbusch, Fischer, and Samuelson (1977) to incorporate an arbitrary number of countries separated by geographic barriers.

3.1 A Ricardian Framework

We describe a world of N countries in which each country can produce every good along the interval $[0, 1]$. Efficiency in the production for any good varies across countries. Aside from these Hicks' neutral efficiency differences, production of any good anywhere combines inputs in the same way, although the price of the bundle of inputs may vary across countries.

Within a country there are many potential producers of any good, but only the most efficient ones will ever be in business. Efficient producers of good j in country i can convert one bundle of inputs into a quantity $Z_{1i}(j)$ of good j (where the subscript 1 indicates that Z_1 corresponds to the *most* efficient method for making j in i).

Below we make the input bundle a Cobb-Douglas combination of labor and intermediate inputs, treating the labor share β as independent of j. For now it is more convenient to keep the inputs bundled together, denoting their price index by w_i. It thus costs $w_i/Z_{1i}(j)$ to produce a unit of good j in country i using the cheapest local means.[6]

Goods can be transported between countries, but at a cost. We make the standard iceberg assumption that delivering one unit of a good in country n requires shipping $d_{ni} \geq 1$ units, and we normalize $d_{ii} = 1$ for all i. We impose the plausible "triangle inequality" on the geographic barrier parameters d_{ni}:

$$d_{ni} \leq d_{nk} d_{ki} \quad \forall k, \tag{1}$$

i.e., an upper bound on the cost of moving goods from i to n is the cost of moving them via some third country k.[7]

Taking into account geographic barriers, the lowest cost in country n of obtaining good j from country i is:

$$C_{1ni}(j) = \left(\frac{w_i}{Z_{1i}(j)} \right) d_{ni} \tag{2}$$

while the lowest cost irrespective of source is:

$$C_{1n}(j) = \min_i \{ C_{1ni}(j) \}. \tag{3}$$

To extract any implications for trade we need to specify how efficiencies are distributed across countries. The fact that we must take minima across a large number of sources makes the model potentially cumbersome. To cut through this problem EK take the distribution of

[6]By assuming that w_i is the same for any good in a country, we ignore possible factor intensity differences, precluding any Heckscher-Ohlin explanation for trade patterns. We could (and intend to) generalize our framework to incorporate differences in factor intensities and in factor endowments as a determinant of specialization, but our goal here is to see how far we get with a purely Ricardian theory.

[7]Arbitrage across markets should ensure that this restriction holds.